

1

## VIRTUAL MACHINE IMAGE ACCESS DE-DUPLICATION

### CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 13/588,050, filed Aug. 17, 2012, and incorporated by reference herein.

### FIELD OF THE INVENTION

Embodiments of the invention generally relate to information technology, and, more particularly, to caching for virtual machine (VM) images.

### BACKGROUND

Challenges exist in the de-duplication of virtual machine (VM) images. Input/output (I/O) contention occurs when accessing VM images, and managing such contention can be important for application performance, user experience, infrastructure cost, etc. However, challenges can arise due to, for example, a high density virtualized environment, a large number of VM images, and/or limited resources (memory and disk space) for caching images locally on compute nodes.

Different VM images often have common portions of data. Reasons for similarity can include, for example, similar operating systems, similar applications, and/or the fact that many new images are created by slightly modifying existing images. Accordingly, VM image access de-duplication aims to avoid I/O operations on blocks with identical content.

Existing approaches include on-demand streaming of a VM image, which includes copy-on-read (CoR), copy-on-write (CoW), and adaptive pre-fetching. Such approaches, however, do not exploit image similarity. Existing approaches can also include the use of a de-duplicated VM image repository. Such approaches attempt to exploit image similarity to combat image sprawl, but lack run-time support (that is, retrieving an image requires reconstituting and copying the entire image).

Other approaches include a general de-duplicated file system, which attempts to exploit file content similarity to reduce disk space occupation, but requires replacing existing file systems. Also, such approaches only consider de-duplicating block allocation instead of file access. Additionally, existing approaches can include VM memory page/cache sharing. Such approaches attempt to discover and share identical memory pages by content scanning or exchanging page information, but introduce high overhead costs.

### SUMMARY

In one aspect of the present invention, techniques for virtual machine image access de-duplication are provided. An exemplary computer-implemented method for de-duplicating virtual machine image accesses can include steps of identifying one or more identical blocks in two or more images in a virtual machine image repository, generating a block map for mapping different blocks with identical content into a same block, deploying a virtual machine image by reconstituting an image from the block map and fetching any unique blocks remotely on-demand, and de-duplicating virtual machine image accesses by storing the deployed virtual machine image in a local disk cache.

Another aspect of the invention or elements thereof can be implemented in the form of an article of manufacture tangibly

2

embodying computer readable instructions which, when implemented, cause a computer to carry out a plurality of method steps, as described herein. Furthermore, another aspect of the invention or elements thereof can be implemented in the form of an apparatus including a memory and at least one processor that is coupled to the memory and operative to perform noted method steps. Yet further, another aspect of the invention or elements thereof can be implemented in the form of means for carrying out the method steps described herein, or elements thereof; the means can include (i) hardware module(s), (ii) software module(s), or (iii) a combination of hardware and software modules; any of (i)-(iii) implement the specific techniques set forth herein, and the software modules are stored in a tangible computer-readable storage medium (or multiple such media).

These and other objects, features and advantages of the present invention will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram illustrating the concept of discovering identical blocks, according to an embodiment of the present invention;

FIG. 2 is a diagram illustrating the concept of image access de-duplication, according to an embodiment of the present invention;

FIG. 3 is a block diagram illustrating an example embodiment, according to an aspect of the invention;

FIG. 4 is a flow diagram illustrating techniques for de-duplicating virtual machine image accesses, according to an embodiment of the invention; and

FIG. 5 is a system diagram of an exemplary computer system on which at least one embodiment of the invention can be implemented.

### DETAILED DESCRIPTION

As described herein, an aspect of the present invention includes content-aware caching for virtual machine images. As noted above, challenges exist in efficient de-duplication of virtual machine images. Existing approaches are based on fixed or variable size chunks, which require significant changes to existing operating systems. Aspects of the invention, as detailed herein, are based on pair-wise similarity between virtual machine images.

At least one embodiment of the invention includes redirecting I/O accesses to an image A to another image B which is available in a memory or local disk cache. Meanwhile, the underlying storage of virtual machine images is in the original format. Accordingly, aspects of the invention include virtual machine image access de-duplication.

As described herein, many accesses to a local disk and/or remote VM repository are for blocks with identical content. One technique for identifying or determining block content includes comparing two or more blocks bit by bit. Another technique includes calculating the hash value of each block and combining all blocks with the same hash value. Accordingly, VM image accesses can be de-duplicated by sharing cache. At least one embodiment of the invention includes translating accesses to different blocks with the same content into accesses to the same blocks. After detecting that a first block and a second block have the same content, upon receipt of a request for the first block, at least one embodiment of the invention includes redirecting the request to the second block.